CrossMark
← click for updates

COMMENTARY

Advancing the activity cliff concept [v1; ref status: indexed, http://f1000r.es/1wf]

Ye Hu, Dagmar Stumpfe, Jürgen Bajorath

Department of Life Science Informatics, B-IT, LIMES Program Unit Chemical Biology and Medicinal Chemistry, Rheinische Friedrich-Wilhelms-Universität, Bonn, D-53113, Germany

Abstract

The activity cliff concept has experienced increasing interest in medicinal chemistry and chemoinformatics. Activity cliffs have originally been defined as pairs of structurally similar compounds that are active against the same target but have a large difference in potency. Activity cliffs are relevant for structure-activity relationship (SAR) analysis and compound optimization because small chemical modifications can be deduced from cliffs that result in large-magnitude changes in potency. In addition to studying activity cliffs on the basis of individual compounds series, they can be systematically identified through mining of compound activity data. This commentary aims to provide a concise yet detailed picture of our current understanding of activity cliffs. It is also meant to introduce the further refined activity cliff concept to a general audience in drug development.

Article Status Summary

Referee Responses

Referees	1	2	3
v1 published 30 Sep 2013	☑ report	☑ report	☑ report

1 **Gerhard Müller**, Mercachem Netherlands

2 **Patrick Walters**, Vertex Pharmaceuticals, Inc. USA

3 **Meir Glick**, Novartis Institutes for BioMedical Research USA

Latest Comments

No Comments Yet

Corresponding author: Jürgen Bajorath (bajorath@bit.uni-bonn.de)

Competing Interests: No competing interests declared.

Introduction

Activity cliffs have been discussed in the medicinal and computational chemistry literature since the early 1990s[1–4]. In the practice of medicinal chemistry, structurally similar compounds with large potency differences are often encountered, mostly during the chemical optimization of individual compound series. Moreover, activity cliffs have also been systematically identified by computational compound data mining[3,4], which has sparked further interest in the activity cliff concept. Hence, in addition to the study of activity cliffs on a case-by-case basis in medicinal chemistry, a large knowledge base of activity cliff information is emerging. In addition, in recent years, the activity landscape concept has also become popular[5]. Activity landscapes are generally defined as graphical representations that integrate similarity and potency relationships of compounds sharing the same biological activity[5], and activity cliffs are their most prominent features[3,5]. As compound data sets rapidly grow in size, activity landscape representations are increasingly used as tools for SAR visualization[6], which further emphasizes the notion of activity cliffs. Hence, the activity landscape and cliff concepts go hand in hand.

Two recent perspective articles have provided detailed accounts of activity cliff research and new developments[3,4]. For an in-depth review of the activity cliff research area, the interested reader is referred to these publications and the references therein. This commentary does not aim to present a full account of activity cliffs and their utility in drug discovery. Rather, it aims to distill out the information that is most relevant to provide a differentiated and critical, yet easy-to-understand view of activity cliffs. In addition, some new findings are reported concerning the target distribution of activity cliffs and coordination of cliffs, which further complement the picture. As mentioned above, computational approaches have substantially influenced our current understanding of activity cliffs. Since we strive for a widely accessible presentation of the activity cliff concept, the discussion of computational details is kept herein to an essential minimum. Furthermore, given that a number of recent activity cliff investigations have originated from our laboratory, some of the views and recommendations presented herein are at least partly subjective. However, it is hoped that they might, nevertheless, stimulate further exploration and discussion of the activity cliff concept. Several recommendations made should also aid in practical applications.

Definition-related key aspects

An activity cliff has originally been defined as a *pair of structurally similar compounds with a large difference in potency*[2,3]. This general definition has four key aspects, which require further consideration and specification:

(i) Only a pair of compounds is considered.

(ii) Both compounds are active (against the same target).

(iii) A *structural similarity criterion* must be specified (i.e., how is *similarity* assessed and how *similar* must compounds be?).

(iv) A *potency difference criterion* must be established (i.e., when is a *potency difference* considered to be *large*?).

In the following, these key points will be further evaluated (in reverse order).

Potency difference criterion

To clearly establish the potency difference criterion, it must not only be decided how large a potency difference between two compounds should be but also considered which type of potency measurements should be utilized. We emphasize that activity cliff information is only useful if the description of cliffs is accurate and interpretable (*vide infra*). This also relates to potency comparisons. Different types of potency measurements should not be combined, e.g., assay-dependent IC_{50} measurements and (in theory) assay-independent equilibrium constants (K_i values) should be separately considered. Moreover, the use of approximate potency annotations (such as "% inhibition") should be avoided to ensure that SAR information encoded by activity cliffs is accurate. As our understanding of activity cliffs has evolved over the years, we have become increasingly conservative in the assessment of cliffs ("conservatism" will indeed be a recurrent theme in our discussion). Therefore, we generally prefer K_i values (*vide infra*), which are in principle the most accurate measurements.

Concerning the magnitude of potency differences, there is no generally applicable rule for the definition of activity cliffs. We have found that statistical significance assessment typically yields data set-dependent results. On the basis of large-scale SAR exploration of many different data sets, we have also concluded that the presence of an at least 100-fold difference in potency as a cliff criterion generally leads to the identification of "significant" activity cliffs in compound data sets from which SAR information can often be deduced. Clearly, this represents a heuristic and not a rule written in stone.

Similarity criterion

Without doubt, the assessment of compound similarity is the most difficult task for activity cliff definition and analysis, for several reasons. The quantification of compound similarity is strongly dependent on chosen molecular representations (descriptors). In addition, there are no generally accepted similarity measures. For activity cliff definition, the calculation of Tanimoto similarity[7] on the basis of different fingerprint representations[7–9] has thus far been most popular[3]. Fingerprints are generally defined as bit representations of molecular structure and/or properties. As such, they are fairly abstract descriptions of compounds. Two fingerprints of different design that have often been used for the description of activity cliffs are the "molecular access system (MACCS) structural keys" (http://accelrys.com)[8], one of the "classical" fingerprints, and the "extended connectivity fingerprint with bond diameter 4 (ECFP4)"[9], a more recent design. MACCS consists of a set of 166 defined structural fragments whose presence or absence in a compound is monitored and ECFP4 is a topological fingerprint that generates varying numbers of atom environments for test compounds. These fingerprints are calculated from molecular graphs and are thus 2D representations. Furthermore, for the purpose of our discussion, it is sufficient to appreciate that the Tanimoto coefficient (Tc) is a similarity measure that ranges from 0 to 1 and quantifies fingerprint overlap as a measure of molecular similarity (i.e., a Tc value of 0 is produced by fingerprints that share no features and a value of 1 by identical fingerprints). A MACCS Tc value of 0.85 (corresponding to Tanimoto similarity of 85%) has often been applied as a similarity criterion for activity cliff formation[3]. This value approx.

corresponds to an ECFP4 Tc value of 0.56 because the same percentage of compound pairs reaches or exceeds these MACCS- and ECFP4-dependent values in systematic compound comparisons[4]. Because computational similarity methods are strongly representation- and compound class-dependent, the assessment of activity cliffs suffers from the same dependencies. Consequently, activity cliff distributions often vary significantly dependent on the representations and similarity measures used[10,11]. Furthermore, calculated fingerprint Tc values are often difficult to interpret from a medicinal chemistry point of view[3–5], which further complicates matters.

In light of these difficulties, attempts have been made to replace calculated similarity values for activity cliff assessment by structurally more conservative and intuitive similarity criteria. For example, a substructure-based similarity criterion has been introduced on the basis of the matched molecular pair (MMP) formalism[12,13]. An MMP is defined as a pair of compounds that are only distinguished by a structural change at a single site[12], i.e., the exchange of a substructure, which is termed a chemical transformation[13]. Importantly, the presence of a defined substructure relationship such as the formation of an MMP can also be applied as a similarity criterion. For the definition of activity cliffs, transformation size-restricted MMPs

have been introduced in which transformations are limited to relatively small and chemically meaningful replacements[14]. Figure 1A and 1B show exemplary Tanimoto similarity- and MMP-based activity cliffs, respectively. The latter activity cliff category has been termed MMP cliff[14]. The similarity criterion underlying MMP cliffs is simple and intuitive. MMP cliffs are often found to further improve the chemical interpretability of activity cliffs compared to cliffs defined on the basis of calculated similarity values[4,14].

Another intuitive categorization of activity cliffs has been introduced on the basis of consistently defined scaffolds (obtained from compounds by removal of R-groups)[15] and the presence of different scaffold/R-group relationships, as illustrated in Figure 1C[16]. This categorization makes it possible, for example, to distinguish activity cliffs that are caused by R-group replacements, small chemical changes in core structures, or chiral centers[16].

Moreover, activity cliffs can also be defined by comparing compound binding modes in complex ligand/target X-ray structures and calculating their 3D similarity[17,18], as shown in Figure 1D. These "3D cliffs" enable the interpretation of activity cliff formation on the basis of experimentally observed ligand-target interactions and
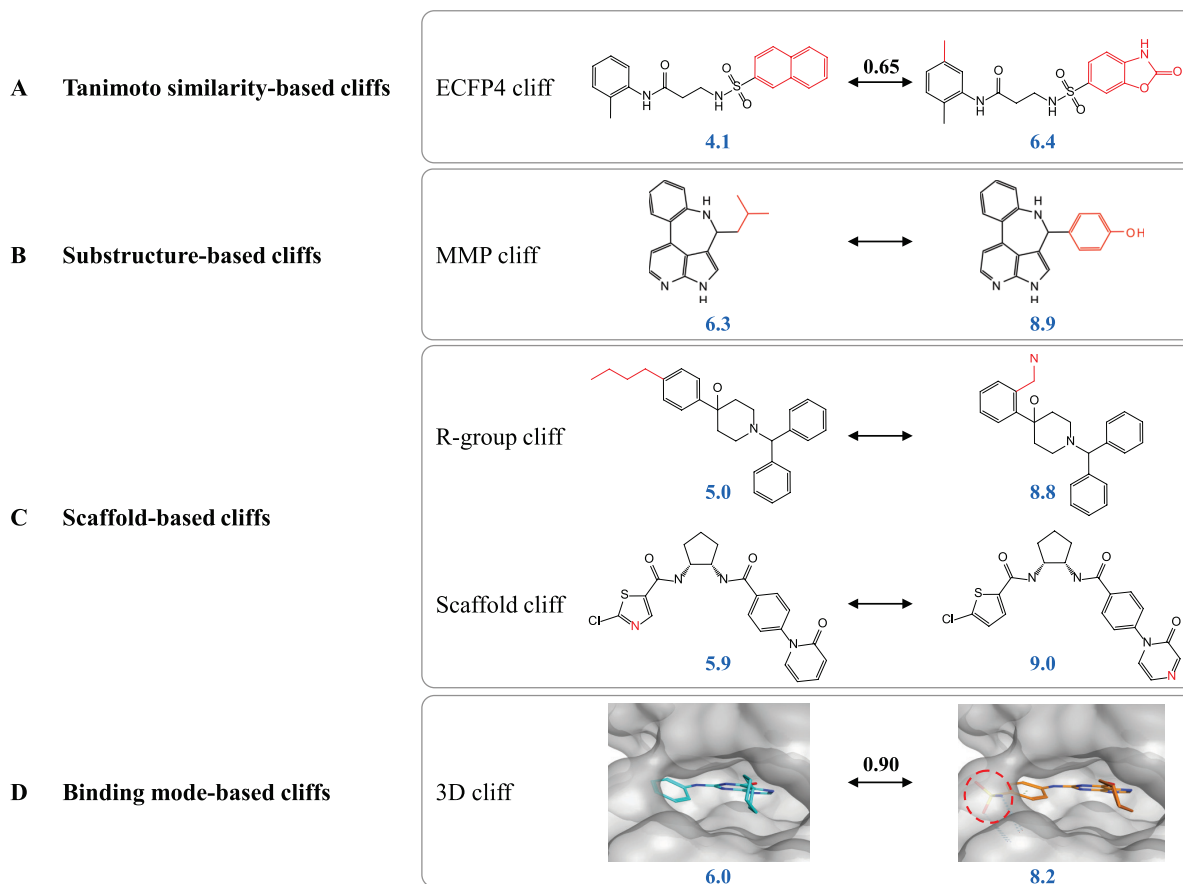
Figure 1. Categorization of activity cliffs. In (**A**)–(**D**), four categories of activity cliffs are shown. Structural differences between activity cliff compounds are highlighted (red). For ECFP4 (**A**) and 3D cliffs (**D**), calculated similarity values are reported, i.e., 0.65 and 0.90 refer to the value of the Tanimoto and property density function similarity coefficient[17], respectively (both of which range from 0 to 1). Compound potency (pK$_i$) values are given in blue.

substantially add to the ligand-centric view of 2D activity cliffs. Although the Protein Data Bank[19] provides a steadily growing source of public domain 3D structures of therapeutically relevant proteins, only small numbers of activity cliffs can be rationalized in three dimensions and compared to 2D cliffs. In addition, 3D cliffs also rely on the calculation of similarity values. The quantification of the 3D similarity of compound binding modes is a fairly complex task because positional and conformational changes need to be taken into account.

Taken together, the examples in Figure 1 illustrate that activity cliffs can be defined in rather different -and more or less intuitive- ways, depending on the applied similarity criteria.

Single-target activity

Based on the original definition of activity cliffs, both compounds are required to be active against a specific target (*vide supra*). Several extensions of the activity cliff concept have been introduced that depart from this theme (for example, by considering selectivity against a pair of targets instead of single-target activity)[3]. In principle, there is no requirement to exclusively consider active compounds for activity cliff assessment. Rather, active and inactive compounds might also be compared, provided confirmed inactive compounds are available for a given target[20]. For SAR analysis in medicinal chemistry, the identification of small structural changes that render compounds active or inactive is of high interest, and the inclusion of confirmed inactive compounds further increases the frequency of activity cliff formation and hence our knowledge base[20]. However, if inactive compounds are taken into consideration, a potency difference criterion is no longer applicable to define activity cliffs. Instead, a potency threshold must be set for active compounds as a cliff criterion. For example, one might require an active compound to have a potency of at least 100 nM to qualify for the formation of an activity cliff with an inactive one. The choice of this threshold is essentially subjective and it might be adjusted, depending on the application. It should also be noted that only small numbers of confirmed inactive compounds are typically available from compound optimization projects. Rather, confirmed inactive compounds mostly result from biological screening campaigns. Thus, care must be taken to obtain high-confidence activity data. For example, confirmatory bioassays from PubChem[21] presently provide a source of confirmed inactive compounds for more than 100 different targets.

Isolated versus coordinated activity cliffs

The definition of activity cliffs on the basis of compound pairs might imply that cliffs are mostly formed in an "isolated" manner. This means that cliff partners are only involved in a single activity cliff and have no structural neighbors with large potency differences. However, this is clearly not the case. For example, series of highly and lowly potent structural analogs have been identified in a variety of compound data sets that form multiple and overlapping activity cliffs[22], giving rise to the notion of "coordinated" activity cliffs[23]. Figure 2 shows an example of a compound set in which highly coordinated activity cliffs are formed. Higher-order activity cliff configurations involving multiple compounds are of particular interest for medicinal chemistry, given their high SAR information content. Such activity cliff arrangements can be systematically

Mitogen-activated protein kinase 14
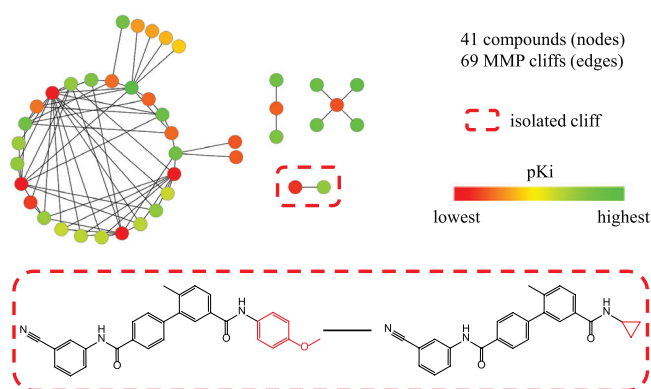
Figure 2. **Isolated versus coordinated cliffs.** MMP cliffs involving inhibitors of mitogen-activated protein kinase 14 are shown in a network representation. Nodes represent compounds that are connected by edges if they form an MMP cliff. Furthermore, nodes are color-coded according to pK$_i$ values of all inhibitors applying a continuous color spectrum from red (lowest potency) via yellow (medium) to green (highest). Only one isolated activity cliff was detected in the inhibitor data set (indicated by a dashed red box). All other activity cliffs were formed in a coordinated manner involving multiple active compounds. Structures of the inhibitors forming the isolated activity cliff are shown.

explored through data mining[22,23]. On the basis of our most recent survey (*vide infra*), only small proportions of activity cliffs are formed in isolation.

Preferred definition

As discussed above, there are many different ways to represent activity cliffs. Is it, then, possible to formulate a generally preferred detailed definition? The answer is: in principle, no; in practice, yes. From first principles, one is unable to determine (at least at present) how compound similarity should best be accounted for. In addition, setting the potency difference criterion for meaningful activity cliff and SAR analysis is subject to heuristic approximations. However, on the basis of our experience with a variety of data analyses and practical applications, we generally prefer the following activity cliff definition[10]:

(a) Similarity criterion: Formation of a size-restricted MMP[14].

(b) Potency difference criterion: At least two orders of magnitude.

(c) Activity measurements: Equilibrium constants.

This MMP cliff definition is conservative both from a compound similarity and activity data perspective and favors chemical interpretability of activity cliffs.

Frequency of occurrence

How often are activity cliffs found in bioactive compounds? What is the proportion of active compounds that participate in the formation of cliffs? Up-to-date results providing answers to these and other questions are reported in Table 1. These results were obtained from a large-scale analysis of compound data sets extracted from ChEMBL

Table 1. Activity cliff frequency.

Types	Percentage (%)			
	Activity cliffs	Cliff-forming compounds	Isolated cliffs	Coordinated cliffs
MACCS	6.8	41.0	1.4	98.6
ECFP4	5.5	35.3	2.2	97.8
MMP cliffs	5.2	27.6	3.5	96.5
3D cliffs	8.5	13.4	20.4	79.6

The average frequency of occurrence of activity cliffs and cliff-forming compounds is reported for 129 target sets[4] and different molecular representations including the MACCS and ECFP4 fingerprints and MMPs. The proportion of activity cliffs was calculated on a per-target basis relative to the total number of compound pairs meeting similarity criteria and the proportion of cliff-forming compounds relative to all active compounds. Also reported is the propensity of 3D activity cliffs based on comparison of ligand binding modes in complex X-ray structures. In each case, isolated and coordinated activity cliffs are distinguished. Potency difference criterion: at least 100-fold on the basis of equilibrium constants. Similarity criteria: For MACCS and ECFP4, Tc values of at least 0.85 and 0.56, respectively[4]; for 3D cliffs, a binding mode similarity coefficient of at least 0.80[17]; for MMP cliffs, formation of a transformation size-restricted MMP[14]. Target sets: A target set is defined as a set of compounds with activity against the same target. Target sets were extracted from ChEMBL[24] on the basis of two selection criteria: Each set had to contain at least 100 compounds and for all compounds equilibrium constants had to be available. The 129 target sets included more than 35,000 unique compounds[4].

(https://www.ebi.ac.uk/chembl/)[24]. Further details are provided in the legend of Table 1. For this survey, equilibrium constants were exclusively used and a potency difference of at least two orders of magnitude was required. Thus, the MMP cliffs reported in Table 1 correspond to our preferred activity cliff definition (*vide supra*). It should be noted that 3D cliffs are statistically underrepresented compared to 2D activity cliffs and that their frequencies of occurrence should not be directly compared. Depending on the molecular representations used, between 5.2% and 6.8% of all qualifying pairs of similar compounds form activity cliffs. The molecular representation dependence of activity cliff assessment is also reflected by the percentage of compounds that participate in activity cliffs, which ranges from 27.6% for MMP- over 35.3% for ECFP4- to 41% for MACCS-based cliffs. MMP cliffs occur with slightly lower frequency than fingerprint-based cliffs and involve a smaller proportion of active compounds. Nevertheless, on average, MMP cliffs are formed by on average every fourth active compound across many different data sets. Hence, even on the basis of this conservative assessment, *activity cliffs frequently occur and provide direct access to SAR information.* We also note that ~96.5%–98.6% of all 2D cliffs are not formed in isolation but in a coordinated manner involving more than two compounds. The example in Figure 2 illustrates that the degree of coordination is often high. For public domain 3D cliffs, the rate of isolated activity cliffs is much higher (20.4%) than for 2D cliffs. This is likely the case because structurally distinct ligands are often chosen for crystallization in order to explore different compound binding modes.

Utilization

Given the considerable frequency with which activity cliffs are formed in different compound sets, a key question is to what extent activity cliff information might currently be utilized in the practice of medicinal chemistry? This question is very difficult to answer since it is hardly possible to systematically track this information with medicinal chemists on a per-project basis. However, data mining studies can provide at least some evidence for the potential utilization of activity cliffs. In a recent study, activity cliffs were systematically identified in compound data sets evolving over time[25]. For each highly potent activity cliff partner, it was determined whether structural analogs of this compound were reported after the activity cliff became available. If structural analogs of a highly potent cliff partner were detected in subsequent years, the possibility existed that activity cliff information provided a starting point for further compound optimization. Alternatively, if no such analogs were identified, no evidence existed for activity cliff progression, as outlined in Figure 3. On the basis of this analysis, evidence for the utilization of activity cliffs was only available for 25% of all available cliffs. By contrast, no evidence for cliff progression was detected for the remaining 75% of activity cliffs[25]. Thus, in light of these findings, *we would conclude that existing activity cliff information is currently under-utilized in the practice of medicinal chemistry*. It is apparently difficult to bridge between data mining investigations and practical medicinal chemistry applications and consider information from the public domain early in the course of compound optimization projects. This would recommend striving for much closer links between chemoinformatics and practical medicinal chemistry.

Another related question should also be of interest. If activity cliff information is utilized, is there an "SAR evolution advantage" detectable compared to other optimization efforts not involving activity cliff compounds? To answer this question, the computational compound pathway model schematically shown in Figure 3 was applied to monitor the progression of activity cliffs towards the most potent compounds in a data set and compare activity cliff-dependent and -independent pathways[26]. Therefore, for each active compound, a search for series of pairwise similar compounds (MMP sequences)
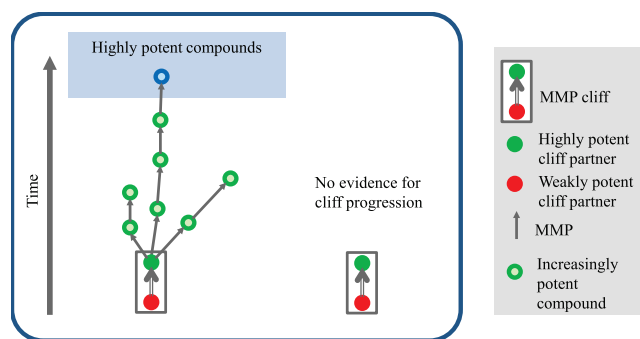
Figure 3. Progression and utilization of activity cliffs. Progression of activity cliffs over time is evaluated by searching for analogs of highly potent activity cliff partners. One of the compound pathways originating from the activity cliff on the left is leading over a sequence of analogs to one of the most potent data set compounds. For the activity cliff on the right, no analogs of the highly potent cliff partner are detected. Thus, in this case, there is no evidence for the potential utilization of activity cliff information.

was carried out that ultimately reached one of the 10% most potent compounds in a data set, thus delineating putative compound optimization paths in accord with the pathway model in Figure 3. Although the mean potency of activity cliff compounds and other active compounds was overall very similar, compound pathways originating from 54% of all activity cliffs successfully reached highly potent compounds, compared to only 28% of pathways originating from compounds not involved in cliff formation[26]. *Hence, activity cliff-dependent pathways reached highly potent compounds with higher frequency than cliff-independent pathways, indicating the presence of activity cliff-associated SAR advantages.*

Target distribution

Are activity cliffs differently distributed in compounds active against different targets? This is another question of considerable interest for medicinal chemistry, which has only recently been addressed[4]. Figure 4 reports the distribution of the frequency of occurrence of MMP cliffs and, in addition, compounds participating in cliff formation for more than 200 different target sets of increasing size (each target set consists of compounds active against a specific

target). The proportion of MMP cliffs relative to all MMPs and the percentage of activity cliff compounds among all active compounds were monitored. In small compound sets, significant frequency fluctuations were observed, as one would expect (for statistical reasons). By contrast, the distribution of activity cliffs and cliff-forming compounds was relatively stable for target sets containing 200 or more compounds. The box plot representations in Figure 5 indicate that *there is surprisingly little variation in the frequency of activity cliffs across many different targets for data sets of moderate to large size* (despite the presence of many different specific ligand-target interactions and binding constraints).

Conclusions

Herein, we have focused on the activity cliff concept and provided a further refined view of activity cliffs. For medicinal chemistry, activity cliffs are of particular interest because they are associated with high SAR information content. For a meaningful assessment of activity cliffs, similarity and potency difference criteria need to be clearly defined. Care must be taken to utilize high-confidence activity data for activity cliff analysis. However, similarity assessment is
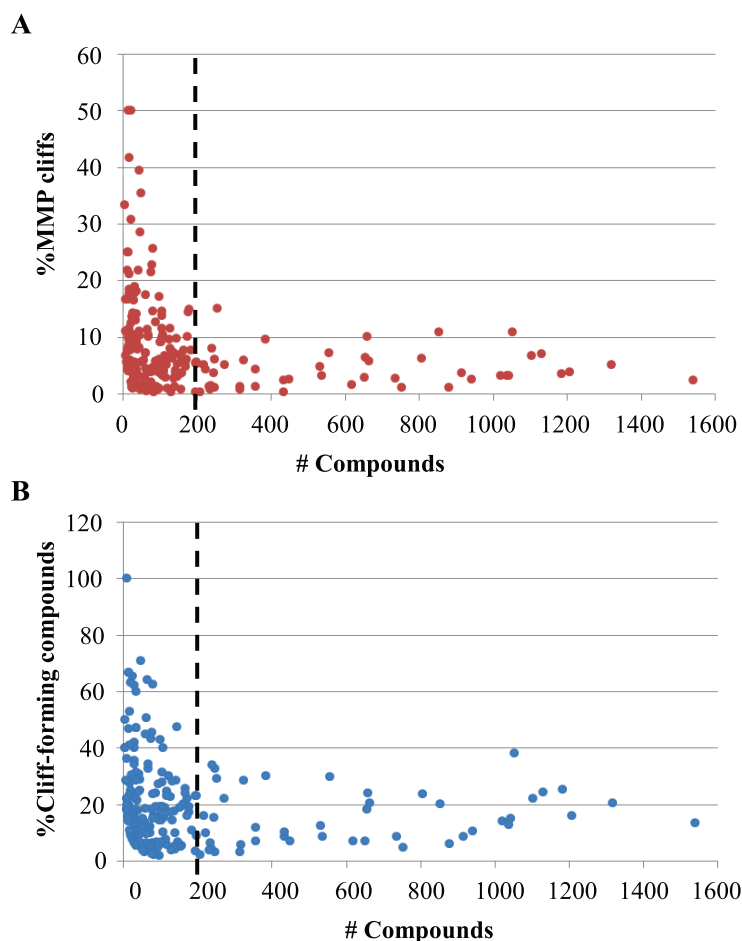
Figure 4. MMP cliffs in target sets of increasing size. In (**A**) and (**B**), the size of target sets is compared to the proportion of MMP cliffs and cliff-forming compounds they contain, respectively. Dots represent individual target sets. Dashed vertical lines mark a target set size of 200 compounds. The analysis was based on 218 target sets with available equilibrium constants extracted from ChEMBL[24].
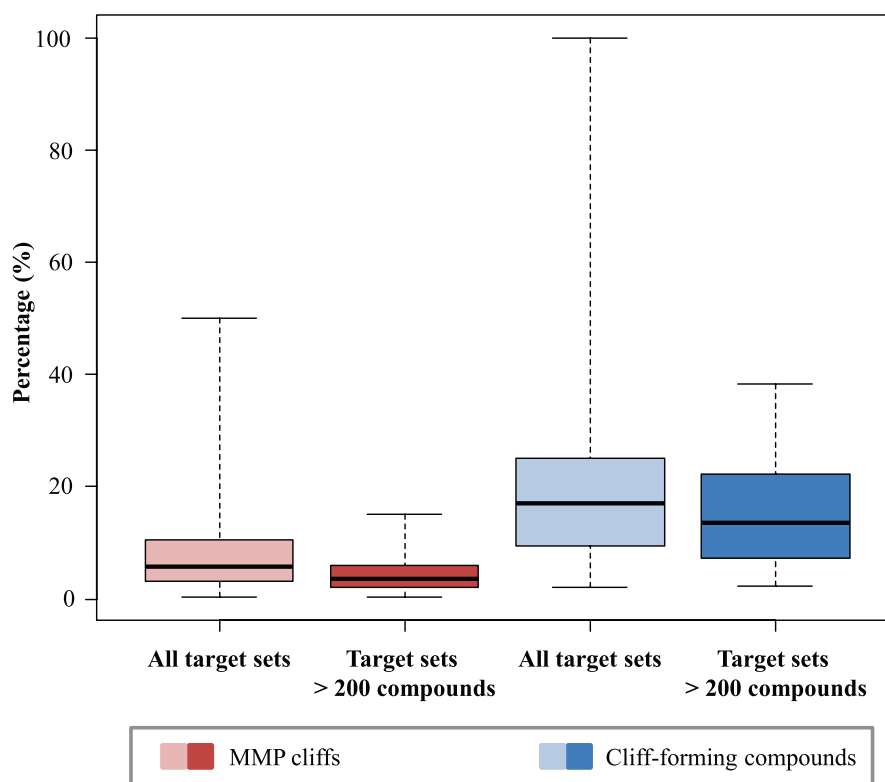
Figure 5. Distribution of MMP cliffs over target sets. The distributions of MMP cliffs (red) and cliff-forming compounds (blue) over target sets are reported as box plots. Each box plot provides the smallest value (bottom line), lower quartile (lower boundary of the box), median value (thick line), upper quartile (upper boundary of the box), and the largest value (top line). The dashed line indicates the value range.

the most critical step in activity cliff analysis. Although calculated similarity values can be conveniently used to describe activity cliffs, they often limit the interpretability of activity cliffs in medicinal chemistry. Therefore, substructure-based activity cliff definitions have been introduced such as MMP cliffs that further support chemical interpretation. Activity cliffs are formed with relatively high frequency among active compounds, indicating that they provide a substantial source of SAR information. In fact, if activity cliff information is utilized, as assessed on the basis of pathway modeling, compound paths originating from activity cliffs more frequently yield highly potent compounds than optimization paths originating from other active compounds. However, there currently is no evidence for utilization of about three quarters of activity cliffs in compound data sets evolving over time, which indicates that available activity cliff information is under-utilized in the practice of medicinal chemistry. Thus, there should be significant potential for further improvement of compound optimization efforts by taking activity cliff information from data mining into consideration. It has also been determined that activity cliffs are relatively evenly distributed across compounds active against a variety of targets,

perhaps surprisingly so. In summary, the activity cliff concept provides an intuitive access to SAR information and can be evaluated from different perspectives. Recent analyses have yielded in part unexpected results that further differentiate our current view of activity cliffs and associated SAR features.

References

1. Lajiness M: **Evaluation of the Performance of Dissimilarity Selection Methodology.** In *QSAR: Rational Approaches to the Design of Bioactive Compounds*; Silipo, C., Vittoria, A., Eds.; Elsevier: Amsterdam, Netherlands, 1991; pp 201–204.

2. Maggiora GM: **On Outliers and Activity Cliffs – Why QSAR often Disappoints.** *J Chem Inf Model.* 2006; **46**(4): 1535.
    PubMed Abstract | **Publisher Full Text**

3. Stumpfe D, Bajorath J: **Exploring Activity Cliffs in Medicinal Chemistry.** *J Med Chem.* 2012; **55**(7): 2932–2942.
    PubMed Abstract | **Publisher Full Text**

4. Stumpfe D, Hu Y, Dimova D, *et al.*: **Recent Progress in Understanding Activity Cliffs and their Utility in Medicinal Chemistry.** *J Med Chem.* 2013; in press.
    PubMed Abstract | **Publisher Full Text**

5. Wassermann AM, Wawer M, Bajorath J: **Activity Landscape Representations for Structure-Activity Relationship Analysis.** *J Med Chem.* 2010; **53**(23): 8209–8223.
    PubMed Abstract | **Publisher Full Text**

6. Stumpfe D, Bajorath J: **Methods for SAR Visualization.** *RSC Adv.* 2012; **2**(2): 369–378.
    Publisher Full Text

7. Willett P, Barnard JM, Downs GM: **Chemical Similarity Searching.** *J Chem Inf Comput Sci.* 1998; **38**(6): 983–996.
    Publisher Full Text

8. **MACCS Structural Keys.** Accelrys: San Diego, CA.

9. Rogers D, Hahn M: **Extended-Connectivity Fingerprints.** *J Chem Inf Model.* 2010; **50**(5): 742–754.
    PubMed Abstract | **Publisher Full Text**

10. Stumpfe D, Bajorath J: **Frequency of Occurrence and Potency Range Distribution of Activity Cliffs in Bioactive Compounds.** *J Chem Inf Model.* 2012; **52**(9): 2348–2353.
    PubMed Abstract | **Publisher Full Text**

11. Medina-Franco JL, Martínez-Mayorga K, Bender A, *et al.*: **Characterization of Activity Landscapes using 2D and 3D Similarity Methods: Consensus Activity Cliffs.** *J Chem Inf Model.* 2009; **49**(2): 477–491.
    PubMed Abstract | **Publisher Full Text**

12. Kenny PW, Sadowski J: **Structure Modification in Chemical Databases.** In *Chemoinformatics in Drug Discovery*; Oprea, T. I., Ed.; Wiley-VCH: Weinheim, Germany, 2005; pp 271–285.
    Publisher Full Text

13. Hussain J, Rea C: **Computationally Efficient Algorithm to Identify Matched Molecular Pairs (MMPs) in Large Data Sets.** *J Chem Inf Model.* 2010; **50**(3): 339–348.
    PubMed Abstract | **Publisher Full Text**

14. Hu X, Hu Y, Vogt M, *et al.*: **MMP-Cliffs: Systematic Identification of Activity Cliffs on the Basis of Matched Molecular Pairs.** *J Chem Inf Model.* 2012; **52**(5): 1138–1145.
    PubMed Abstract | **Publisher Full Text**

15. Bemis GW, Murcko MA: **The Properties of Known Drugs. 1. Molecular Frameworks.** *J Med Chem.* 1996; **39**(15): 2887–2893.
    PubMed Abstract | **Publisher Full Text**

16. Hu Y, Bajorath J: **Extending the Activity Cliff Concept: Structural Categorization of Activity Cliffs and Systematic Identification of Different Types of Cliffs in the ChEMBL Database.** *J Chem Inf Model.* 2012; **52**(7): 1806–1811.
    PubMed Abstract | **Publisher Full Text**

17. Hu Y, Bajorath J: **Exploration of 3D Activity Cliffs on the Basis of Compound Binding Modes and Comparison of 2D and 3D cliffs.** *J Chem Inf Model.* 2012; **52**(3): 670–677.
    PubMed Abstract | **Publisher Full Text**

18. Hu Y, Furtmann N, Gütschow M, *et al.*: **Systematic Identification and Classification of Three-dimensional Activity Cliffs.** *J Chem Inf Model.* 2012; **52**(6): 1490–1498.
    PubMed Abstract | **Publisher Full Text**

19. Berman HM, Westbrook J, Feng Z, *et al.*: **The Protein Data Bank.** *Nucleic Acids Res.* 2000; **28**(1): 235–242.
    PubMed Abstract | **Publisher Full Text** | **Free Full Text**

20. Hu Y, Maggiora GM, Bajorath J: **Activity cliffs in PubChem confirmatory bioassays taking inactive compounds into account.** *J Comput Aided Mol Des.* 2013; **27**(2): 115–124.
    PubMed Abstract | **Publisher Full Text**

21. Wang Y, Xiao J, Suzek TO, *et al.*: **PubChem's BioAssay Database.** *Nucleic Acids Res.* 2012; **40**(Database issue): D400–D412.
    PubMed Abstract | **Publisher Full Text** | **Free Full Text**

22. Vogt M, Huang Y, Bajorath J: **From Activity Cliffs to Activity Ridges: Informative Data Structures for SAR Analysis.** *J Chem Inf Model.* 2011; **51**(8): 1848–1856.
    PubMed Abstract | **Publisher Full Text**

23. Namasivayam V, Bajorath J: **Searching for Coordinated Activity Cliffs Using Particle Swarm Optimization.** *J Chem Inf Model.* 2012; **52**(4): 927–934.
    PubMed Abstract | **Publisher Full Text**

24. Gaulton A, Bellis LJ, Bento AP, *et al.*: **ChEMBL: A Large-scale Bioactivity Database for Drug Discovery.** *Nucleic Acids Res.* 2012; **40**(Database issue): D1100–D1107.
    PubMed Abstract | **Publisher Full Text** | **Free Full Text**

25. Dimova D, Heikamp K, Stumpfe D, *et al.*: **Do Medicinal Chemists Learn from Activity Cliffs? A Systematic Evaluation of Cliff Progression in Evolving Compound Data Sets.** *J Med Chem.* 2013; **56**(8): 3339–3345.
    PubMed Abstract | **Publisher Full Text**

26. Stumpfe D, Dimova D, Heikamp K, *et al.*: C**ompound Pathway Model to Capture SAR progression: Comparison of Activity Cliff-Dependent and -Independent Pathways.** *J Chem Inf Model.* 2013; **53**(5): 1067–1072.
    PubMed Abstract | **Publisher Full Text**

Current Referee Status:

Referee Responses for Version 1

Meir Glick
Center for Proteomic Chemistry, Novartis Institutes for BioMedical Research, Cambridge, MA, USA

Approved: 16 October 2013

Referee Report: 16 October 2013
The commentary provides an up to date view of activity cliffs and assess their prevalence in ChEMBL. In addition, it discusses how activity cliffs can be defined using potency and similarity criteria. The paper is well written.

My only suggestion is to make table 1 more detailed; that is to include the activity cliff frequency around target families or even around the 129 individual target sets.

I have read this submission. I believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Competing Interests: No competing interests were disclosed.

Patrick Walters
Vertex Pharmaceuticals, Inc., Cambridge, MA, USA

Approved: 09 October 2013

Referee Report: 09 October 2013
This paper provides an overview of the concept of activity cliffs in structure-activity relationships and provides a number of examples of how the technique can be applied in a drug discovery setting. Activity cliffs can provide a means of better understanding the relationships between chemical structure and biological activity, as well as identifying regions of a molecule that are critical to activity. The review is directed toward the non-expert and should provide an adequate background for those unfamiliar with the topic. Those familiar with activity cliffs may want to consult more specialized reviews written by the same authors. This review is well-written, easy to follow, and provides many key references for those who would like to learn more about the topic.

The authors' breakdown of the activity cliff concept into components will be helpful for those unfamiliar with the topic. It may be useful to point out that utilizing log differences in activity can make activity cliff scores more intuitive to chemists and biologists. Those with a limited knowledge of chemoinformatics will appreciate the discussion of fingerprint similarity. However, a figure might make these concepts even more clear. The authors' point out that "*calculated fingerprint Tc values are often difficult to interpret from a medicinal chemistry point of view*". In my experience, Tanimoto similarity typically makes sense, this sentence might be better as "*calculated fingerprint Tc values are *sometimes* difficult to interpret*". The authors mention matched molecular pairs (MMP) as a means of expressing molecular similarity.

In addition to references 12 and 13, it would be useful to include what I consider to be the seminal paper on matched molecular pairs.
Leach, A. G.; Jones, H. D.; Cosgrove, D. A.; Kenny, P. W.; Ruston,L.; MacFaul, P.; Wood, M.; Colclough, N.; Law, L. Matched Molecular Pairs as a Guide in the Optimization of Pharmaceutical Properties; A Study of Aqueous Solubility, Plasma Protein Binding and Oral Exposure. J. Med. Chem. 2006, 49, 6672–6682.

As chemists and biologists pursue more complex targets, it is necessary to increase the sophistication of our SAR analysis. I appreciate the fact that the authors point out how activity cliffs can also be used to examine multi-target activity relationships and links between sometimes-distant functional groups.

I agree with the authors' assessment that the activity cliff concept is underutilized in medicinal chemistry. Hopefully with the introduction of robust commercial software tools for the identification and exploitation of activity cliffs, the technique will be more widely used.

This review provides and excellent introduction to activity cliffs for those unfamiliar with the method. The paper is well referenced and will hopefully inspire a new group of users to continue to advance the method.

I have read this submission. I believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Competing Interests: No competing interests were disclosed.

Gerhard Müller
Medicinal Chemistry, Mercachem, Nijmegen, Netherlands

Approved: 07 October 2013

Referee Report: 07 October 2013
Ye Hu, Dagmar Stumpfe, and Jürgen Bajorath elaborate in their commentary entitled '***Advancing the activity cliff concept'*** on the application of an abstract chemoinformatics-based analysis of differentiated structure-activity relationships within given compound sets. Given the high level of complexity and the expert computational chemistry knowledge often required to comprehend and to apply those latest developments in the field of chemoinformatics, I consider this commentary as extremely helpful in bridging exactly that obvious gap between the practitioners of bioactive compound data mining on the one hand, and the laboratory-based practitioners of medicinal chemistry on the other hand, respectively.

Within the context of today's practice of medicinal chemistry, the sound and conclusive interpretation of structure-activity data is one of the most relevant, though challenging assignments in any small molecule lead finding and optimization campaign. In the course of those projects, involved teams are permanently required to make high-impact decisions on e.g. further prosecution of specific chemotypes based on systems that are, from an experimental point of view, notoriously under determined. Hence, it is of great value to account for all publicly available information on the compound-centric aspects in a given drug discovery project in order to make the most educated decision, specifically on the 'tractability' of a given chemotype.

More specifically, compound classes that after a certain chemical elaboration fail to exhibit a differentiated

activity profile, i.e. do not show activity cliffs, should not be prosecuted any further since all too often those 'flat SAR' or 'rolling-hill' analogues were repeatedly identified as cul-de-sac series in many lead finding campaigns. In that context, the authors provide very helpful guidelines for the medicinal chemistry community on how to access the concept of activity cliffs by sharing expert insights they have gained over many systematic studies on practical aspects such as choosing the right potency difference and molecular similarity criteria to reliably define activity cliff patterns. Different molecular similarity metrics are introduced and comparatively ranked on their applicability and interpretability within a potential medicinal chemistry project. We all might have heard about Tanimoto, MACCS and ECFP4 fingerprints; but I assume that not too many of us are acquainted to the MMP (matched molecular pair) concept and utilize that concept pro-actively in pursued lead finding and optimization campaigns.

The group around Jürgen Bajorath actively de-mystifies those similarity concepts for us by providing quantitative measures on the inter-relation, thereby sharing their experience on the practicability of all those computational-based approaches. And for a medicinal chemist it is inspiring to see how e.g. scaffold-based cliffs unfold in a comprehensible way. I find it especially useful that the authors share their preferred definitions and settings in terms of applied similarity criteria, potency difference criteria, as well as activity measurements, thus highlighting scope and limitations of the underlying methodology. The more interested reader who wants to get really deeply involved into algorithmic details or further applications of the activity cliff concept is referred to the most recent and relevant research references.

Within the framework of a baseline study on 35000 bioactive compounds acting at 129 targets in the ChEMBL database, the overall occurrence of activity cliffs is analysed underlining the finding that activity cliffs occur quite frequently and contain valuable intrinsic SAR information that is not applied systematically. It becomes apparent that activity cliff-dependent optimization pathways seem to have a cliff-associated SAR advantage over other optimization strategies. Based on those findings, I can only support the claim that is made by the authors, notably to establish much closer links between chemoinformatics and practising medicinal chemists for a better and earlier integration of existing activity cliff information into pursued projects.

Despite the multiple and rather different ways towards activity cliffs, the concept as such becomes a tangible and intuitive asset that should be applied more frequently in the today's pre-clinical discovery arena.

I see especially the attempt of reaching highest user-friendliness for a putative end-user that renders this commentary as very helpful to get the medicinal chemistry community alerted for this useful but still under-appreciated concept. The group around Jürgen Bajorath undoubtedly qualifies as an advocate in that sense, and the community of practising medicinal chemists should start to move in their direction accordingly.

I have read this submission. I believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Competing Interests: No competing interests were disclosed.